# 30,000 Mailbox Resiliency Solution for Microsoft Exchange 2013 using Lenovo Converged HX7510 Appliances and Hyper-V

Last update: 30 August 2016

**Provides a technical overview of Lenovo Converged HX Solution for Microsoft Exchange 2013**

**Contains performance results using Microsoft ESRP Storage Program**

**Uses HX7510 appliances with Microsoft 2012 R2 Hyper-V hypervisor**

**Shows class-leading hyper-converged performance of HX7510 appliance**

**Pawan Sharma**

**Mike Perks**

LENOVO PRESS

# Table of Contents

# 1 Overview

This document provides information on Lenovo Converged HX7510 solution for Microsoft Exchange Server, based the Microsoft Exchange *Solution Reviewed Program (ESRP) – Storage* program[1].

The intended audience for this document is technical IT architects, system administrators, and managers who are interested in executing virtualized Microsoft Exchange workloads on the Lenovo Converged HX Series Nutanix Appliances. The abbreviated term of HX Series is used in the remainder of this document.

HX Series appliances provide a hyper-converged infrastructure which seamlessly pools compute and storage to deliver high performance for the virtual workloads and provides flexibility to combine the local storage using a distributed file system to eliminate shared storage such as SAN or NAS. These factors make the solution cost effective without compromising the performance.

This document describes the solution for executing 30,000 mailboxes with Microsoft Exchange Server 2013 on Lenovo Converged HX7510 appliances using the Nutanix Acropolis Hypervisor (AHV). The mailbox resiliency data group with two database copies executes on 2 four node clusters.

The cluster of 4 Lenovo Converged HX7510 appliances proved more than capable of handling the high IOPs generated by 30,000 mailboxes. Part of the reason for this is because each Lenovo Converged HX7510 appliance uses 3 HBAs for the 24 drives. As a consequence the cluster has an average IOPS rate that is 23% better and an average latency improvement of 2 milliseconds over other systems with less HBAs.

Chapter 2 provides a technical overview of the HX Series appliances and explains why the combination of Lenovo servers and Nutanix software provides best of breed system performance and reliability.

Chapter 3 provides a solution overview of Microsoft Exchange on Lenovo Converged HX7510 appliances with the deployment best practices, recommended appliance configuration, and the performance test results from the Microsoft ESRP storage test. The detailed log files are in appendix A.

## 1.1 Disclaimer

This document has been produced independently of Microsoft Corporation. Microsoft Corporation expressly disclaims responsibility for, and makes no warranty, express or implied, with respect to, the accuracy of the contents of this document.

The information contained in this document represents the current view of Lenovo on the issues discussed as of the date of publication. Due to changing market conditions, it should not be interpreted to be a commitment on the part of Lenovo, and Lenovo cannot guarantee the accuracy of any information presented after the date of publication.

---

[1] The *ESRP – Storage* program was developed by Microsoft Corporation to provide a common storage testing framework for vendors to provide information on its storage solutions for Microsoft Exchange Server software. For more details on the *Microsoft ESRP – Storage* program, use the following website: http://technet.microsoft.com/en-us/exchange/ff182054.aspx.

# 2 Technical overview of HX Series appliances

Lenovo Converged HX Series appliances are designed to help you simplify IT infrastructure, reduce costs, and accelerate time to value. These hyper-converged appliances from Lenovo combine industry-leading hyper-convergence software from Nutanix with Lenovo enterprise platforms.

## 2.1 Server Components

The Lenovo Converged HX series appliances are available in a variety of E5-2600 v4 processor-based models to support different workloads. In particular the Lenovo Converged HX7510 is recommended for Microsoft Exchange and is shown below


The table below provides a summary of the configuration options for the HX7510.

| Appliance Model | Form Factor | Solid State Drives (SSD) | Hard Disk Drives (HDD) | Memory | CPU | Dual port NIC | GPUs |
|---|---|---|---|---|---|---|---|
| HX7510 | 2.5" drives 2U | Standard<br>4 x 480 GB<br>4 x 800 GB<br>4 x 1200 GB<br>4 x 1600 GB<br>Encrypted<br>4 x 400 GB<br>4 x 800 GB<br>4 x 1600 GB | Standard<br>20 x 1TB<br>20 x 2TB<br>Encrypted<br>20 x 2TB | 128 GB<br>256 GB<br>384 GB<br>512 GB<br>768 GB | 2 x 2643 v4<br>2 x 2680 v4<br>2 x 2697 v4<br>2 x 2699 v4 | 1 or 2 | N/A |

For more information see the product guide for Lenovo Converged HX7510: lenovopress.com/lp0507

## 2.2 Software components

This section gives an overview of the software components used in the solution.

### 2.2.1 Hypervisor

The HX Series appliances support the following hypervisors:

- Nutanix Acropolis Hypervisor based on KVM (AHV)
- VMware ESXi 5.5 U3B
- VMware ESXi 6.0 U1B
- VMware ESXi 6.0 U2
- Microsoft Windows 2012 R2 DataCenter Edition (Hyper-V)

The HX Series appliances come standard with the AHV preloaded in the factory. The other hypervisors are supported as a field-installable option using the Nutanix Foundation tool.

### 2.2.2 Nutanix Prism

Nutanix Prism gives administrators a simple and elegant way to manage virtual environments. Powered by advanced data analytics and heuristics, Prism simplifies and streamlines common workflows within a data center. Nutanix Prism is a part of the Nutanix software preloaded on the appliances and offers the following features:

- Single point of control
    - Accelerates enterprise-wide deployment
    - Manages capacity centrally
    - Adds nodes in minutes
    - Supports non-disruptive software upgrades with zero downtime
    - Integrates with REST APIs and PowerShell
- Monitoring and alerting
    - Tracks infrastructure utilization (storage, processor, memory)
    - Centrally monitors multiple clusters across multiple sites
    - Monitors per virtual machine (VM) performance and resource usage
    - Checks system health
    - Generates alerts and notifications
- Integrated data protection
    - Offers customizable RPO/RTO and retention policies
    - Supports configurable per-VM replication (1:1, 1:many and many:1)
    - Provides efficient VM recovery
    - Deploys affordable data recovery (DR) and backup to the cloud
- Diagnostics and troubleshooting
    - Provides time-based historical views of VM activity
    - Performs proactive alert analysis
    - Correlates alerts and events to quickly diagnose issues
    - Generates actionable alerts and reduces resolution times
    - Analyzes trending patterns for accurate capacity planning

### 2.2.3 Nutanix Controller VM

The Nutanix Controller VM (CVM) is the key to hyper-converged capability and each node in a cluster has its own instance. Figure 1 shows the main components of the CVM.

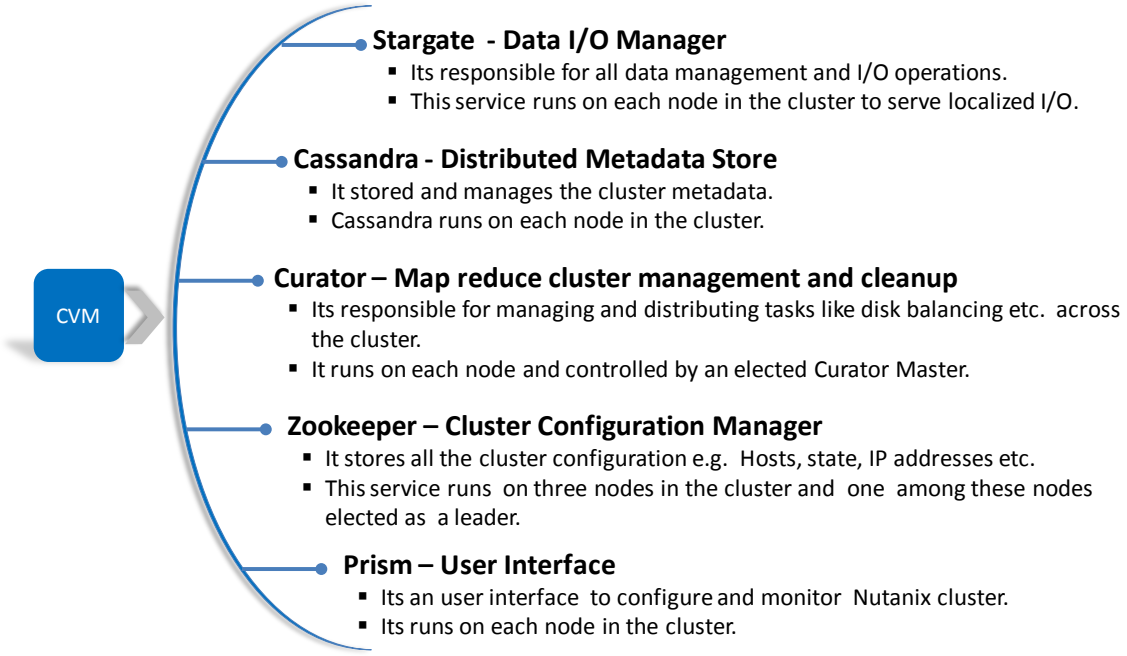


***Figure 1: Controller VM components***

The CVM works as interface between the storage and hypervisor to manage all I/O operations for the hypervisor and user VMs running on the nodes as shown in Figure 2.

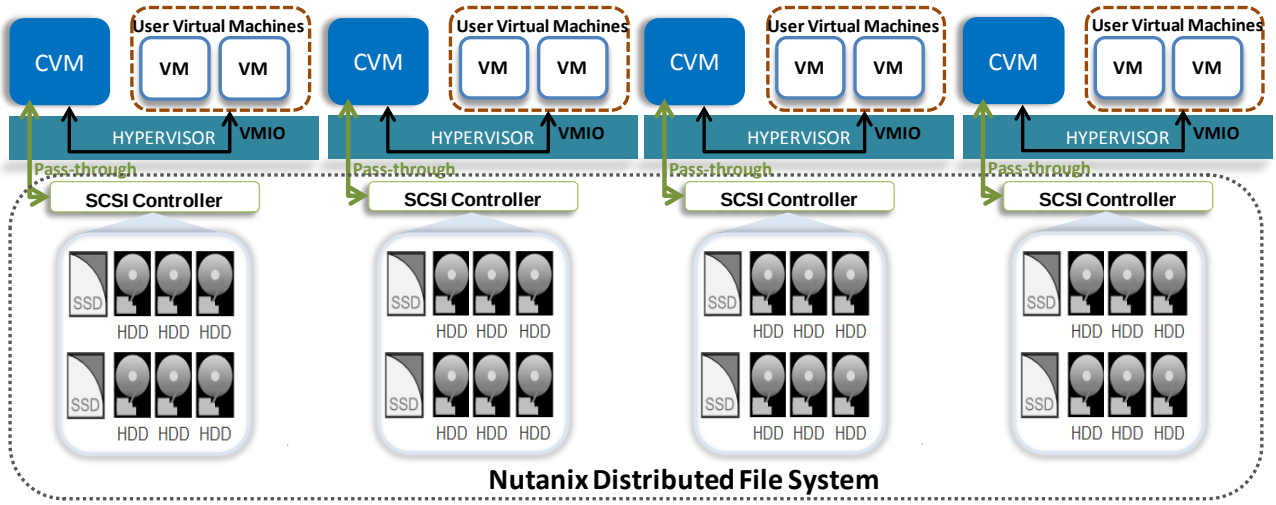


***Figure 2: CVM interaction with Hypervisor and User VMs***

CVM virtualizes all the local storage attached to each node in a cluster and presents it as centralized storage array using Nutanix Distributed File System (NDFS). All I/O operations are handled locally to provide the highest performance. See section 2.4 for more details on the performance features of NDFS.

### 2.2.4 Nutanix Foundation

Nutanix Foundation is a separate utility that you use to orchestrate the installation of hypervisors and Nutanix software on one or more nodes. The maximum number of nodes that can be deployed at one time is 20.

Foundation is available both as a stand-alone VM and also integrated into the CVM. Because CVM is

pre-installed in the factory, the CVM integration of Foundation simplifies the deployment and cluster creation of new servers delivered from the factory.

## 2.3 Networking components

A cluster of HX Series appliances requires both a 10 GbE network for data and a 1GbE network for hardware management.

### 2.3.1 10GbE networking

Each Lenovo Converged HX Series appliance contains 1 or 2 dual-port 10GbE network adapters as well as 4 on-board 1GbE ports. The hypervisors are configured by the Nutanix software so that all of the network ports on the appliance (both 10GbE and 1 GbE) are pooled. The hypervisor VM management network should use the same 10GbE network.

Because all of the network ports are pooled, each appliance only needs two network IP addresses; one for the hypervisor and one for the Nutanix CVM. These IP addresses should be all on the same subnet.

It is recommended that two top of rack (TOR) switches are used for redundancy. The second network adapter provides an additional level of redundancy in case one of the network adapters fails entirely. Redundancy across the two TORs is provided using bonded Inter-Switch-Links (ISL).

Figure 3 shows 4 HX Series appliances each with one dual-port NIC connected to two TOR switches.

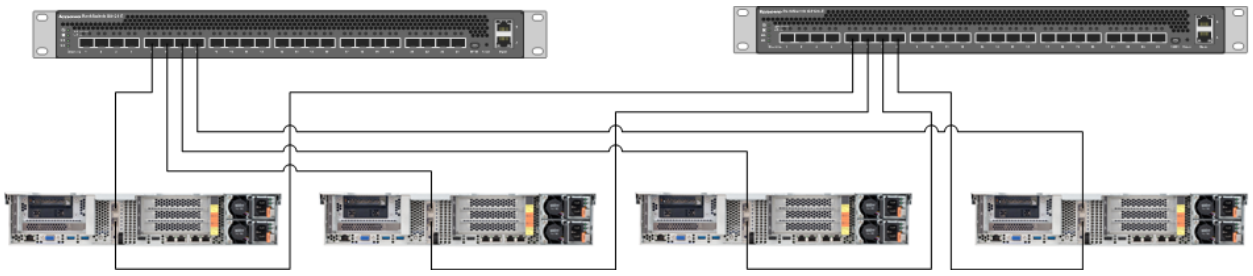

***Figure 3: One NIC and two TOR switches***

Figure 4 shows 4 HX Series appliances each with two dual-port NICs connected to two TOR switches.

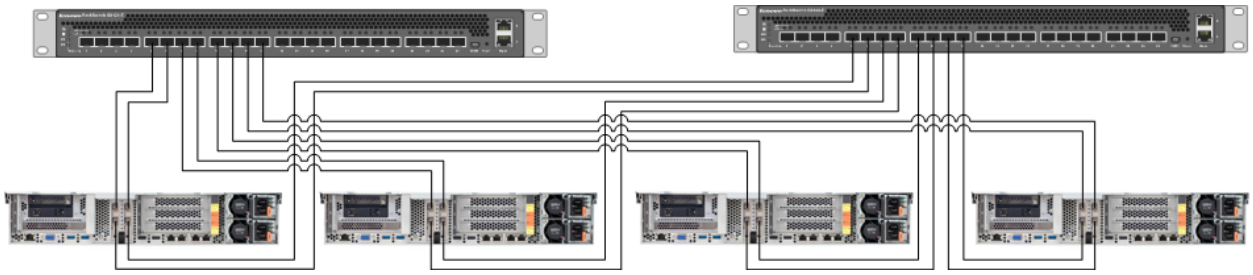

***Figure 4: Two NICs and two TOR switches***

The following Lenovo 10GbE TOR switches are recommended for use in a HX Series cluster:

- Lenovo RackSwitch G8124E
- Lenovo RackSwitch G8272

## Lenovo RackSwitch G8124E

The Lenovo RackSwitch™ G8124E (as shown in Figure 5) is a 10 GbE switch that is specifically designed for the data center and provides a virtualized, cooler, and easier network solution. The G8124E offers 24 10 GbE ports in a 1U footprint. Designed with top performance in mind, the RackSwitch G8124E provides line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data and large data center grade buffers to keep traffic moving.

The G8124E switch is virtualized by providing rack-level virtualization of networking interfaces. The G8124E switch also supports Virtual Fabric, which allows for the distribution of a physical NIC into 2 to 8 vNICs and creates a virtual pipe between the adapter and the switch. The G8124E switch is easier to manage with server-oriented provisioning by using point-and-click management interfaces.


***Figure 5: Lenovo RackSwitch G8124E***

For more information, see this website: lenovopress.com/tips0787

## Lenovo RackSwitch G8272

The Lenovo RackSwitch G8272 uses 10Gb SFP+ and 40Gb QSFP+ Ethernet technology and is specifically designed for the data center. It is an enterprise class Layer 2 and Layer 3 full featured switch that delivers line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data. Large data center-grade buffers help keep traffic moving, while the hot-swap redundant power supplies and fans (along with numerous high-availability features) help provide high availability for business sensitive traffic.

The RackSwitch G8272 (shown in Figure 6), is ideal for latency sensitive applications, such as high-performance computing clusters and financial applications. In addition to the 10 Gb Ethernet (GbE) and 40 GbE connections, the G8272 can use 1 GbE connections.

***Figure 6: Lenovo RackSwitch G8272***

The RackSwitch G8272 supports Lenovo Virtual Fabric, which helps clients significantly reduce cost and complexity that are related to I/O requirements of many virtualization deployments. Virtual Fabric helps reduce the number of multiple I/O adapters to a single dual-port 10 GbE adapter and the number of cables and required upstream switch ports.

By using Virtual Fabric, you can carve a dual-port 10 Gb server adapter into eight virtual network ports (vPorts) and create dedicated virtual pipes between the adapter and switch for optimal performance, higher availability, and improved security. With Virtual Fabric, you can make dynamic changes and allocate bandwidth per vPort so that you can adjust it over time without downtime.

For more information, see this website: lenovopress.com/tips1267

### 2.3.2 1 GbE networking

The dedicated IMM port on all of the Lenovo Converged HX series appliances needs to be connected to a 1GbE TOR switch as shown in Figure 7.

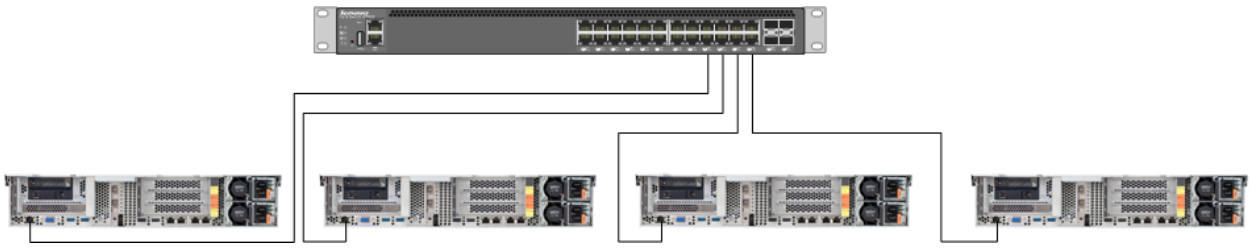

***Figure 7: IMM 1GbE management network***

The following Lenovo 1GbE TOR switches are recommended for use in a HX Series cluster:

- Lenovo RackSwitch G7028
- Lenovo RackSwitch G8052

#### Lenovo RackSwitch G7028

The Lenovo RackSwitch G7028 (as shown in Figure 8) is a 1 Gb top-of-rack switch that delivers line-rate Layer 2 performance at an attractive price. G7028 has 24 10/100/1000BASE-T RJ45 ports and four 10 Gb Ethernet SFP+ ports. It typically uses only 45 W of power, which helps improve energy efficiency.

***Figure 8. Lenovo RackSwitch G7028***

For more information, see this website: lenovopress.com/tips1268.

#### Lenovo RackSwitch G8052

The Lenovo System Networking RackSwitch G8052 (as shown in Figure 9) is an Ethernet switch that is designed for the data center and provides a virtualized, cooler, and simpler network solution. The Lenovo RackSwitch G8052 offers up to 48 1 GbE ports and up to four 10 GbE ports in a 1U footprint. The G8052 switch is always available for business-sensitive traffic by using redundant power supplies, fans, and numerous high-availability features.

***Figure 9: Lenovo RackSwitch G8052***

For more information, see this website: lenovopress.com/tips0813.

### 2.3.3 VLANs

It is a networking best practice to use VLANs to logically separate different kinds of network traffic. The following standard VLANs are recommended:

- Management Used for all management traffic for the hyerpvisor
- Storage network Used for NDFS storage traffic

In addition, each workload application might require one or more VLANs for its logical networks. For larger networks with many workloads, it is easy to run out of unique VLANs. In this case, VXLANs could be used.

The procedure for configuring VLANs in a cluster of HX Series appliances is outside of the scope of this document.

## 2.4 Reliability and performance features

Reliability and excellent performance are important for any workload but particularly for hyper-converged infrastructures like the HX Series appliances.

Reliability, high availability and excellent performance are provided through the following design features of Nutanix software combined with Lenovo Servers.

### Hardware reliability

Lenovo uses the highest quality hardware components combined with firmware that is thoroughly tested. As a consequence Lenovo System x servers have been rated #1 in hardware reliability for the last 2 years. This is important as it lowers the frequency of a server failure which in turn lowers OPEX.

A HX appliance has redundant hardware components by including two power supplies, multiple chassis fans, two Intel CPUs, multiple memory DIMMs, multiple SSDs and HDDs, and optionally up to two dual-port network interface cards.

### Hardware performance

The HX Series appliances have been carefully designed for performance. In addition to all of the usual attributes like processors and memory, the 24 drive HX7510 uses three HBA controllers instead of the one. As a consequence the latency is halved for some workloads that heavily utilize the cold tier. This allows a higher throughput and improved transaction rates.

### Distributed file system

The Nutanix Distributed file system (NDFS) is an intelligent file system which virtualizes the local attached storage (SSD/HDD) on all the nodes in a cluster and presents it as single storage entity to cluster. Figure 10 shows the high level structure of NDFS:

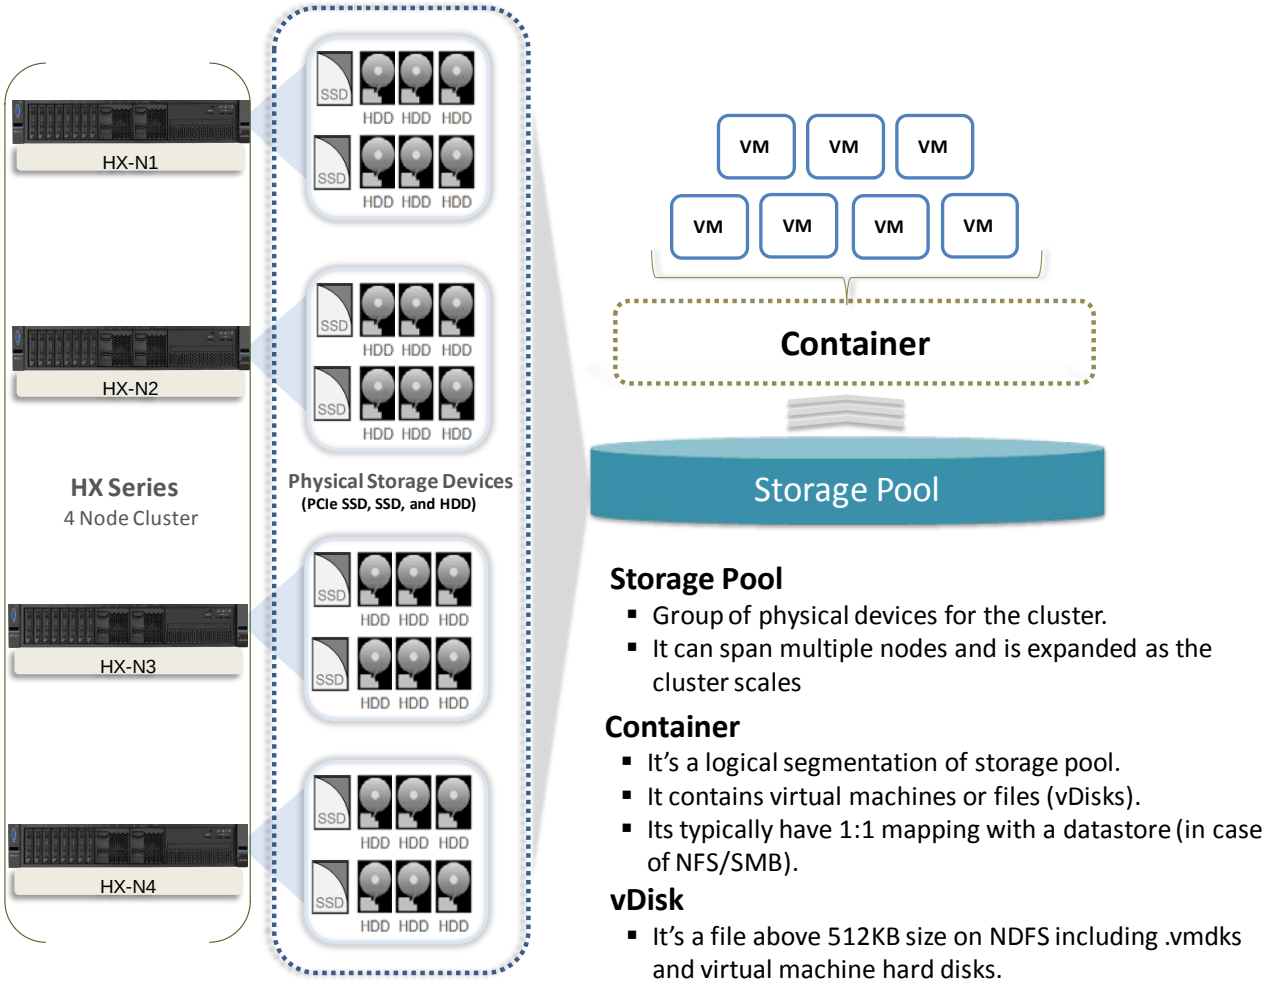


***Figure 10: Nutanix Distributed File System***

## Data protection via replication

The Nutanix platform replication factor (RF) and checksum is used to ensure data redundancy and accessibility in the event of a node or disk failure or corruption. It uses an OpLog which acts as a staging area for incoming writes on low latency SSDs which are then replicated to the OpLogs for one or two other Controller VMs before acknowledging a successful write. This approach ensures that data available in at least two to three different locations and is fault tolerant. While the data is being written a checksum is calculated and stored as part of its metadata.

In the case of a drive or node failure, that data is replicated out to more nodes to maintain the replication factor. A checksum is computed every time the data is read to ensure the data validity. If the checksum and data mismatch, then the data replica is read to replace the invalid copy.

## Performance with data tiering

Nutanix uses a disk tiering concept in which disk resources (SSD and HDD) are pooled together to form a cluster wide storage tier. This tier can be accessed by any node within the cluster for data placement and can leverage the full tier capacity. The following data tiering functions are provided:

- The SSD on a local node always has the highest tier priority for write I/O.
- If the local node's SSD is full then the other SSDs in the cluster are used for I/O.
- The NDFS Information Lifecycle Management (ILM) component migrates cold data from the local SSD to HDD to free up SSD space. It also moves heavily accessed data to the local SSD to provide high performance.

## Performance by data locality

Data locality is a crucial factor for cluster and VM performance. In order to minimize latency the CVM will work to ensure that all I/O happens locally. This ensures optimal performance and provides very low latencies and high data transfer speeds that cannot be achieved easily with shared storage arrays, even if all-flash.

The following occurs in case of a VM migration or high availability event that moves a VM from Node-A to Node-B:

- The VM's data is provided by the CVM running on Node-B.
- All write I/O requests occur locally i.e. to the local storage of Node-B.
- When a request comes for reading old data, the I/O request is forwarded by Node-B to Node-A. NDFS detects that the I/O request originated from different node and migrates the data locally in the background i.e. from Node-A to Node-B so that all subsequent read I/Os are served locally. This approach (migration only on a read) helps to avoid network flooding.

## Performance of snapshots and clones

NDFS provides support for offloaded snapshots and clones using a redirect-on-write algorithm. When a snapshot or clone is created, the base vDisk is marked as read only and another vDisk is created with read/write permissions as shown in Figure 11 and Figure 12 below.

At this point both vDisks have the same block map - a metadata mapping of the vDisk to its corresponding extents. This approach reduces the overhead of creating snapshots and allows snapshots to be taken very quickly with little performance impact.

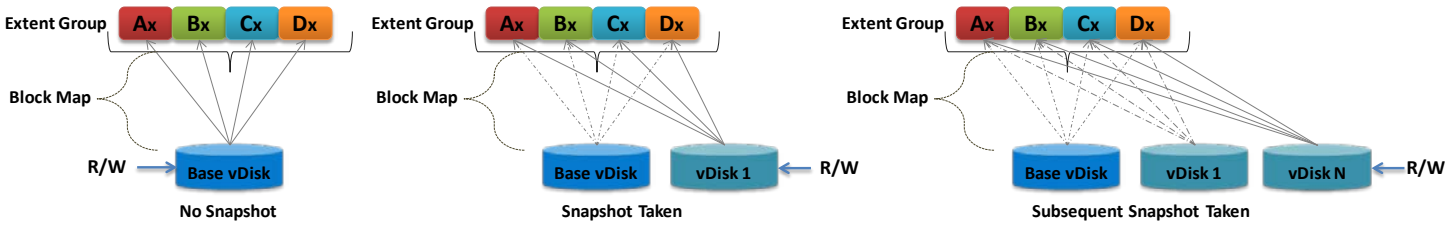


***Figure 11: VM snapshots***

When a VM is cloned the current block map is locked and then clones are created. These updates are metadata only so again no actual I/O takes place. The logic applies for clones of clones as well where a previously cloned VM acts as a base vDisk. All the clones inherit the prior block map and any new writes take place on the individual block maps.

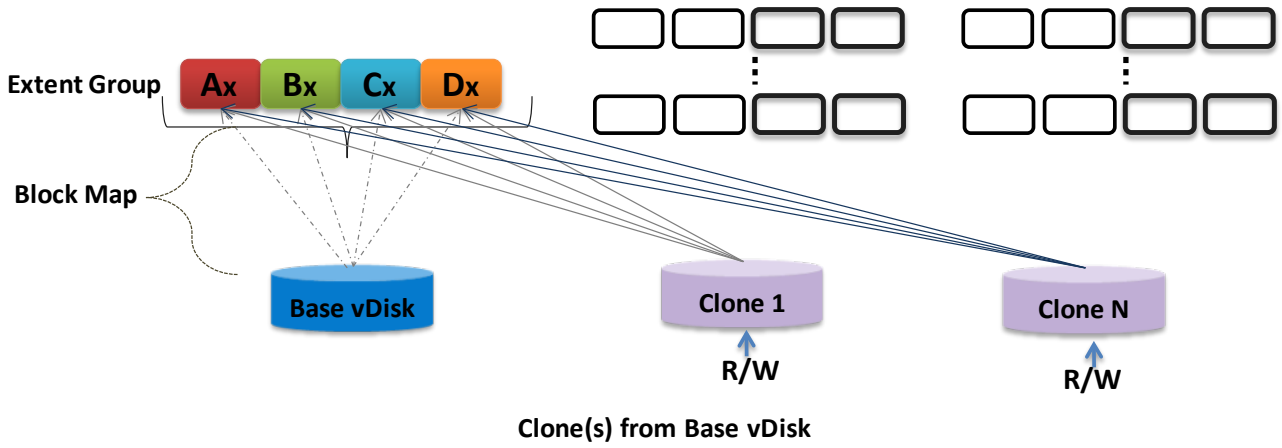


***Figure 12: VM clones***

## Storage reduction via De-duplication and Compression

The Nutanix elastic de-duplication engine increases the effective capacity of a disk, as well as the RAM and cache of the system by removing duplicate data. It's an intelligent technology which performs following actions to increase storage efficiency:

- Sequential streams of data fingerprinted at 4K granularity
- Single instance of the shared VM data is loaded into the cache upon read
- Each node in a cluster performs its own fingerprinting and deduplication

The Nutanix capacity optimization engine is responsible for performing data transformations and compression to achieve data optimization. NDFS provides following compression methods:

- In-line compression sequential streams of data or large I/O sizes are compressed in memory before written to the disk
- Post-process compression whereby data is written in an uncompressed state and the curator framework is used to compress the data in a cluster wide manner

The Nutanix capacity optimization engine uses the Google snappy compression library to deliver good compression ratios with minimal compute overhead and very fast compression or decompression rates.

## Elimination of "split-brain" errors

In a distributed system it is possible for one participant to become disconnected which will cause differences in the stored data. NDFS uses the proven "Paxos" algorithm to eliminate these "split-brain" issues by reaching a consensus (quorum) among the participants in a distributed system before the writes are made.

## Drive reliability via active monitoring

The CVM actively monitors the performance of every drive in a node. The deterioration of a drive's performance may indicate that the drive is about to fail. The CVM proactively moves data off the drive before it fails and marks the drive offline and in need to replacement. The idea is to avoid the expensive data transfers to maintain data redundancy and possible loss of data.

# 3 Microsoft Exchange

Microsoft Exchange Server 2013 is the market leader in enterprise messaging and collaboration. Exchange Server 2013 builds upon the Exchange Server 2010 architecture and was redesigned for simplicity of scale, improved hardware utilization, and increased failure isolation. The goal of Exchange Server 2013 is to support people and organizations as their work habits evolve from a communication focus to a collaboration focus.

## 3.1 Solution overview

Figure 13 shows the architectural overview of the Microsoft Exchange solution using Lenovo Converged HX7510 appliances. This chapter does not address integrating Exchange with unified messaging solutions and handling edge transport routing and distribution.

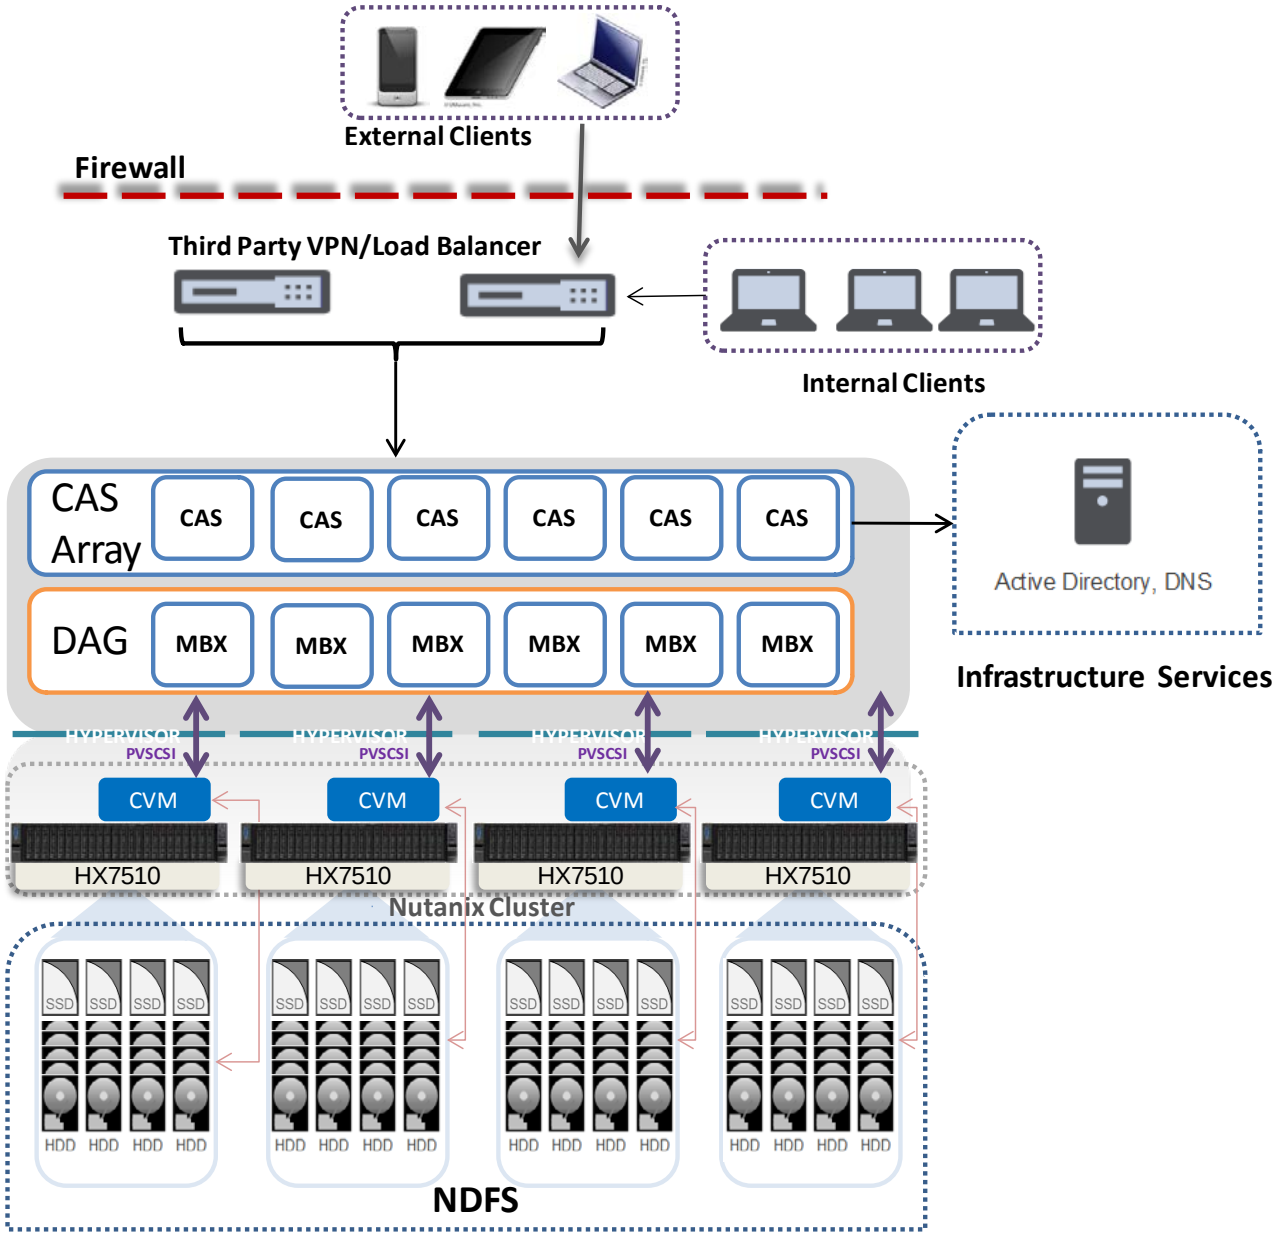


***Figure 13. Lenovo Converged HX Series solution with Microsoft Exchange***

The Client Access Server (CAS) role provides client protocols, SMTP, and unified messaging support. The Mailbox Server (MBX) role provides all of the data processing services. Lenovo recommends that these roles are combined into a multi-role server.

For load balancing into the CAS layer either a network load balancer can be used with a CAS array object or a layer 4 or layer 7 load balancer can be used without the need for configuring a CAS array.

## 3.2 Component model

This section describes the logical component view of the Exchange Server 2013 environment. Figure 14 shows a high-level component model.

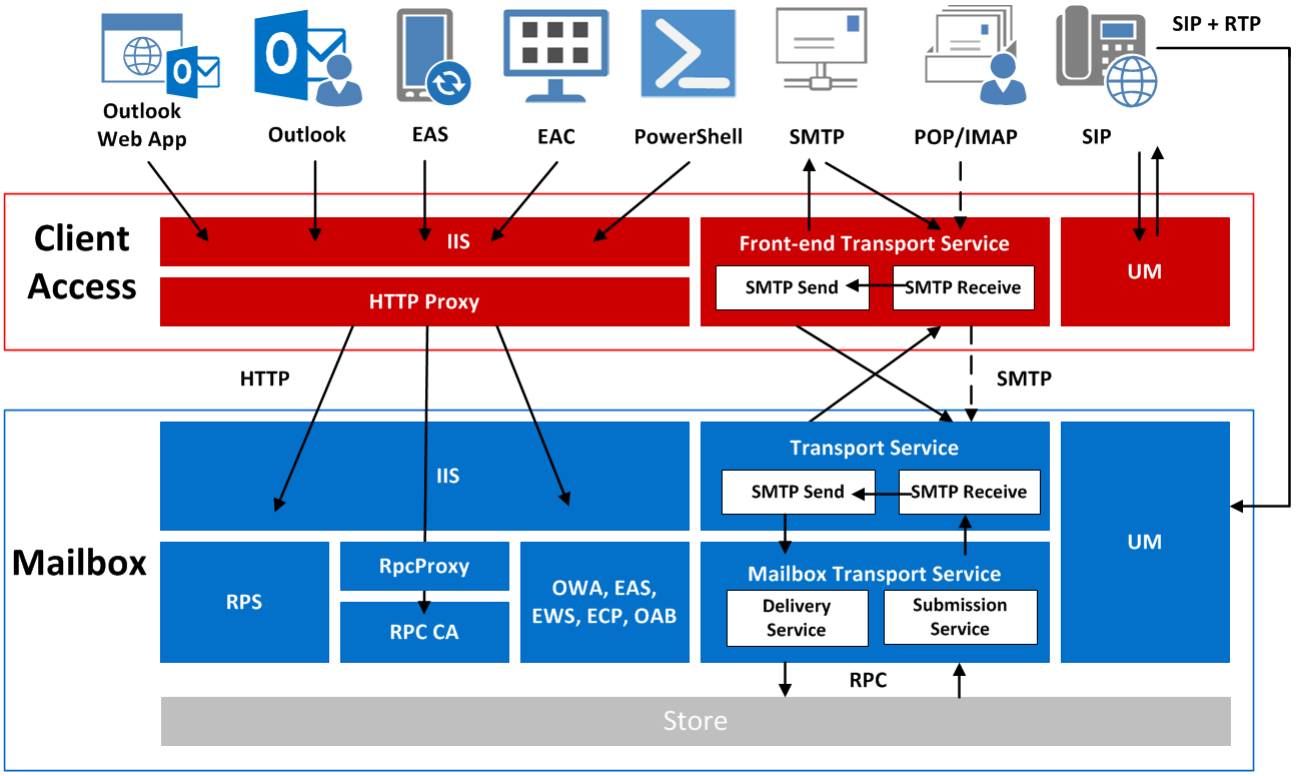


***Figure 14. Exchange Server 2013 logical component view***

The following basic concepts and terminology are used throughout this section:

**Exchange Admin Center (EAC)** – The EAC is the web-based management console in Microsoft Exchange Server 2013 that is optimized for on-premises, online, and hybrid Exchange deployments. The EAC replaces the Exchange Management Console (EMC) and the Exchange Control Panel (ECP), which were the two interfaces used to manage Exchange Server 2010.

**Exchange Control Panel (ECP)** – The ECP is a web application that runs on a Client Access Server and provides services for the Exchange organization.

**Exchange Web Services (EWS)** – EWS provides the functionality to enable client applications to communicate with the Exchange server.

**Internet Information Services (IIS)** – IIS is an extensible web server that was created by Microsoft for use with Windows NT family.

**Internet Message Access Protocol (IMAP**) – IMAP is a communications protocol for email retrieval and storage developed as an alternative to POP.

**Microsoft Exchange ActiveSync (EAS)** – EAS is a communications protocol that is designed for the synchronization of email, contacts, calendar, tasks, and notes from a messaging server to a smartphone or other mobile device.

**Microsoft Outlook® Web App (OWA)** – OWA (formerly Outlook Web Access) is a browser-based email client with which users can access their Microsoft Exchange Server mailbox from almost any web browser.

**Offline Address Book (OAB)** – The OAB is a copy of an address list collection that was downloaded so a Microsoft Outlook user can access the address book while disconnected from the server. Microsoft Exchange generates the new OAB files and then compresses the files and places them on a local share.

**Outlook Anywhere** – Outlook Anywhere is a service that provides RPC/MAPI connectivity for Outlook clients over HTTP or HTTPS by using the Windows RPC over HTTP component. In previous versions of Exchange Server, this function was used for remote or external access only. However, in Exchange Server 2013, all Outlook connectivity is via HTTP/HTTPS (even for internal clients).

**Post Office Protocol (POP)** – The POP is an application-layer Internet standard protocol that is used by local email clients to retrieve email from a remote server over a TCP/IP connection

**Real-time Transport Protocol (RTP)** – RTP is a network protocol for delivering audio and video over IP networks.

**Remote PowerShell (RPS)** – RPS allows you to use Windows PowerShell on your local computer to create a remote Shell session to an Exchange server if you do not have the Exchange management tools installed.

**RPC Client Access (RPC)** – In Microsoft Exchange Server 2007, the Client Access server role was introduced to handle incoming client connections to Exchange mailboxes. Although most types of client connections were made to the Client Access server, Microsoft Office Outlook still connected directly to the Mailbox server when it was running internally with the MAPI protocol.

A new service was introduced with Exchange Server 2010 to allow these MAPI connections to be handled by the Client Access server. The RPC Client Access service provides data access through a single, common path of the Client Access server, with the exception of public folder requests (which are still made directly to the Mailbox server). This change applies business logic to clients more consistently and provides a better client experience when failover occurs.

**Remote Procedure Call over HTTP** – The RPC over HTTP component wraps RPCs in an HTTP layer that allows traffic to traverse network firewalls without requiring RPC ports to be opened. In Exchange 2013, this feature is enabled by default because Exchange 2013 does not allow direct RPC connectivity.

**Session Initiation Protocol (SIP)** – SIP is a protocol that is used for starting, modifying, and ending an interactive user session that involves multimedia elements, such as video, voice, and instant messaging.

**Simple Mail Transfer Protocol (SMTP)** – SMTP is an Internet standard for email transmission.

**Unified Messaging (UM)** – UM allows an Exchange Server mailbox account that was enabled for UM to receive email, voice, and fax messages in the Inbox.

## 3.3 Exchange deployment best practices

This section describes recommended best practices for Microsoft Exchange mailboxes. See also this website for Nutanix Best Practices Guide: Virtualizing Microsoft Exchange: go.nutanix.com/virtualizing-microsoft-exchange-converged-infrastructure.html.

### 3.3.1 Data optimization

By default all Nutanix storage containers are thin provisioned which reduces unused capacity and automatically provisions additional storage capacity when needed. It is also very easy to add additional storage capacity for mailboxes by simply adding nodes to the cluster. It is also possible to set a storage reservation to guarantee a minimum amount of storage capacity.

Data compression can be used to further increase data capacity especially for data that is less frequently accessed. Lenovo recommends enabling compression with a delay of 1440 minutes (1 day) which minimizes the performance impact on I/O writes.

Data de-duplication is not recommended and should be disabled for active Exchange mailboxes because of the frequency of changes. Note that de-duplication may be beneficial for backup volumes which are not changed very often.

A resiliency factor of 2 is the default. This provides a minimum level of data redundancy but a resiliency factor of 3 might be important in some environments. Using erasure coding saves significant storage capacity but it only recommended for archive data.

### 3.3.2 Cluster high availability

The minimum number of nodes in each cluster is 3 and should be at least 4 to provide failover. The following high availability features are recommended for Hyper-V based cluster:

A database availability group (DAG) is the base component of the high availability and site resilience framework that is built into Microsoft Exchange Server 2013. A DAG is a group of up to 16 mailbox servers that hosts a set of mailbox databases and provides automatic database-level recovery from failures that affect individual servers or databases.

A DAG is a boundary for mailbox database replication, database and server switchovers, failovers, and an internal component called *Active Manager*. Active Manager, which runs on every server in a DAG, manages switchovers and failovers.

Any server in a DAG can host a copy of a mailbox database from any other server in the DAG. When a server is added to a DAG, it works with the other servers in the DAG to provide automatic recovery from failures that affect mailbox databases (such as a disk failure or server failure).

Lenovo recommends a DAG configuration of 2 database copies and optionally one lagged copy. With a data resiliency factor of 2, the effective number of copies of each mailbox is 4 and this allows two disk failures without losing data.

DR across datacenters can also be done using DAGs assuming there is sufficient band-width between the sites. The scenarios for active-active and active-passive DR sites using DAGs are outside the scope of this document.

### 3.3.3 Other best practices

Consider the following points regarding virtualizing Exchange:

- All Exchange 2013 server roles should be supported in a single VM.
- Some hypervisors include features for taking snapshots of VMs. VM snapshots capture the state of a VM while it is running. This feature enables you to take multiple snapshots of a VM and then revert the VM to any of the previous states by applying a snapshot to the VM. However, VM snapshots are not application aware, and the use of snapshots can have unintended and unexpected consequences for a server application that maintains state data, such as Exchange. Therefore, making VM snapshots of an Exchange guest VM is not supported.
- Disable Hyper-threading.
- The operating system for an Exchange guest machine must use a disk that has a size equal to at least 15 GB plus the size of the virtual memory that is allocated to the guest machine. This requirement is necessary to account for the operating system and paging file disk requirements. For example, if the guest machine is allocated 16 GB of memory, the minimum disk space that is needed for the guest operating system disk is 31 GB.

## 3.4 DAG architecture overview

The following section illustrates the Lenovo Converged HX7510 based Exchange 2013 mailbox resiliency solution Lenovo implemented for ESRP testing.

**Compute per node:** 2 x Intel E5-2699v4 (22 cores @ 2.2 GHz) processors

**RAM per node:** 512GB

**Raw storage per node:** SSD: 4 x 800GB SATA SSDs – 3.2TB
HDD: 20 x 2TB SATA HDDs – 40TB

**Raw storage per cluster:** SSD: 16 x 800GB SATA SSDs – 12.8TB
HDD: 80 x 2TB SATA HDDs – 160 TB

Figure 15 describes the high level DAG architecture of the 30,000 mailbox virtualized Exchange 2013 mailbox resiliency solution. This solution comprises of two Lenovo Converged HX7510 clusters with 4 nodes per cluster. Each node uses the Microsoft 2012 R2 Hyper-V hypervisor.

The DAG has six Exchange 2013 mailbox servers and two database copies. The two database copies were placed on two physically isolated clusters.

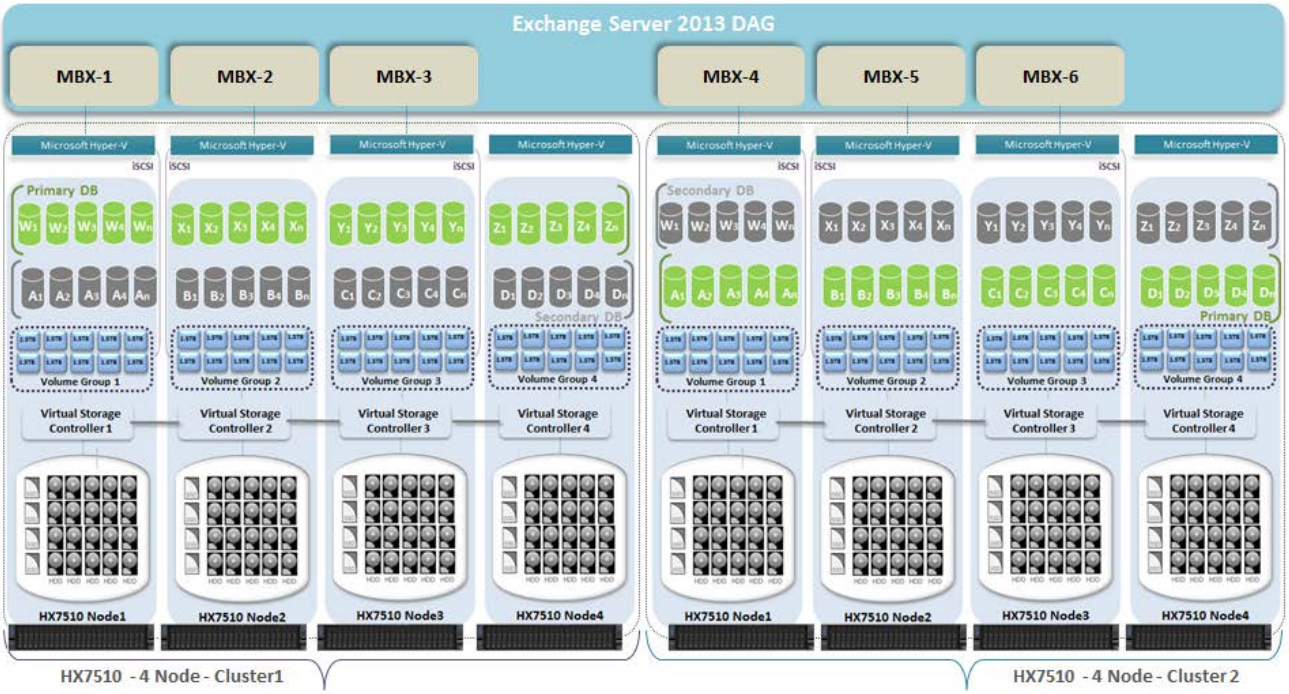


***Figure 15. DAG Architecture using Exchange 2013 and Lenovo Converged HX7510 (30,000 Mailboxes)***

On both Cluster1 and Cluster2, the Mailbox server (MBX) VMs were created on three nodes. All the Database/Logs volumes were connected to Mailbox servers using the NFS protocol. The primary database copy and secondary database copy are stored on two physically separated and isolated clusters. The two clusters can be located at the same datacentre or two different datacenters.

Each Mailbox Server VM is configured as follows:

- 24 vCPUs
- 400 GB RAM
- Windows Server 2012 R2
- Microsoft Exchange 2013 with 10,000 mailboxes

The ESRP-Storage program focuses on storage solution testing to address performance and reliability issues with storage design. However, storage is not the only factor to take into consideration when designing a scale up Exchange solution. Other factors which affect the server scalability are: server processor utilization, server physical and virtual memory limitations, resource requirements for other applications, directory and network service latencies, network infrastructure limitations, replication and recovery requirements, and client usage profiles. All these factors are beyond the scope for ESRP-Storage. Therefore, the number of mailboxes hosted per server as part of the tested configuration may not necessarily be viable for some customer deployment.

For more information on identifying and addressing performance bottlenecks in an Exchange system, please refer to Microsoft's Troubleshooting Microsoft Exchange Server Performance, available at http://technet.microsoft.com/en-us/library/dd335215.aspx.

## 3.5 Targeted customer profile

The target customer profile for a medium enterprise Microsoft Exchange 2013 environment is as follows:

- 30,000 mailboxes of 1GB
- 6x Exchange 2013 servers (3x Tested)
- 0.06 IOPS per mailbox
- 24/7 background database maintenance
- Mailbox resiliency factor of 2
- 10 databases per host

## 3.6 Tested deployment environment

The section describes the tested deployment environment.

### 3.6.1 Simulated exchange configuration

The following table summarizes the simulated Exchange configuration.

| | |
|---|---|
| Number of Exchange mailboxes simulated | 30,000 |
| Number of Database Availability Groups (DAGs) | 1 |
| Number of servers/DAG | 6 (3 tested) |
| Number of active mailboxes/server | 5,000 (5,000 active and 5,000 passive mailboxes per server and tested 10,000 active mailboxes per server) |
| Number of databases/host | 10 |
| Number of copies/database | 2 |
| Number of mailboxes/database | 1000 |
| Simulated profile: I/O's per second per mailbox (IOPS, include 20% headroom) | 0.06 IOPs / Mailbox |
| Database/LUN size | 1.5 TB |
| Total database size for performance testing | 30 TB |
| % storage capacity used by Exchange database[2] | 63.46% |

[2] Storage performance characteristics change based on the percentage utilization of the individual disks. Tests that use a small percentage of the storage (~25%) may exhibit reduced throughput if the storage capacity utilization is significantly increased beyond what is tested in this paper.

### 3.6.2 Storage hardware

The following table summarizes the storage hardware.

| | |
|---|---|
| Storage Connectivity (Fiber Channel, SAS, SATA, iSCSI) | iSCSI |
| Storage model and OS/firmware revision | HX7510 running Microsoft 2012 R2 Hyper-V |
| Storage cache | 76.8GB per node |
| Number of storage controllers | 4x virtual controller virtual machines |
| Number of storage ports | 2x 10 Gbe Port |
| Maximum bandwidth of storage connectivity to host | 20 Gbps per node |
| Switch type/model/firmware revision | Lenovo RackSwitch G8124E (10GbE)<br>Firmware version: 7.7.5 |
| HBA model and firmware | N2215 SAS/SATA HBA |
| Number of HBA's/host | 3 |
| Host server type | 3x Lenovo Converged HX7510 (2 x Intel E5-2699v4 @ 2.2 GHz) 512GB RAM |
| Total number of disks tested in solution | 96 (4 node cluster) |
| Maximum number of spindles can be hosted in the storage | 96 (cluster can be scaled to 40+ nodes) |

### 3.6.3 Storage software

The following table summarizes the storage software.

| | |
|---|---|
| HBA driver | N/A |
| HBA QueueTarget Setting | N/A |
| HBA QueueDepth Setting | N/A |
| Hypervisor | Microsoft 2012 R2 Hyper-V |
| Exchange VM guest OS | Windows Server 2012 R2 |
| ESE.dll file version | 15.00.0847.030 |
| Replication solution name/version | N/A |

### 3.6.4 Storage disk configuration (mailbox store disks)

The following table summarizes the storage disk configuration.

| | |
|---|---|
| Disk type, speed and firmware revision | Per Node:<br>4x Intel S3610 800 GB SATA G3HS 2.5" SSD<br>20x 2 TB 7.2K 6Gbps NL SATA 2.5" G3HS 512e HDD |
| Raw capacity per disk (GB) | 2048 GB (2 TB) |
| Number of physical disks in test | 96 (16x 800 GB + 80x 2 TB) |
| Total raw storage capacity (GB) | 172,800 GB |
| Disk slice size (GB) | N/A |
| Number of slices or disks per LUN | N/A |
| Raid level | Nutanix Replication Factor 2 (RAID 1) |
| Total formatted capacity | 70.82 GB |
| Storage capacity utilization | 40.9% |
| Database capacity utilization | 17.3% |

## 3.7 Performance test results

This section provides a high-level summary of the results of executing the Microsoft ESRP storage test version 4.0 on the configuration of 4 Lenovo Converged HX7510 appliances as described in the previous section. ESRP storage test results include reliability, storage performance, and database backup/restore.

Note that the ESRP program is not designed to be a benchmarking program and tests are not designed to get the maximum throughput for a giving solution. Rather, the program is focused on producing recommendations from vendors for the Exchange application. Therefore, the data presented in this document should not be used for direct comparisons among the solutions and customers should not quote the data directly for their pre-deployment verifications. It is recommended that a proof of concept is carried out to validate the storage design for a specific customer environment.

The results in this section were developed by Lenovo and reviewed by the Microsoft Exchange Product team.

### 3.7.1 Reliability

Several of the tests in the ESP test framework are used to check reliability and run for 24 hours. The test objective is to verify that the storage can handle high I/O workloads for extensive periods. Log and database files are analyzed for integrity after the stress test to ensure there is no database or log corruption.

Executing this test on the Lenovo Converged HX7510 appliances showed:

- No errors reported in the saved event log file.
- No errors reported during the database and log checksum process.

For detailed log results see “24-Hour stress/reliability test results” on page 24.

### 3.7.2 Storage performance results

The primary storage performance test in the ESP test framework is designed to exercise the storage with a maximum sustainable Exchange I/O pattern for 2 hours. The purpose is to reveal how long it takes for the storage to respond to I/O operations under a load.

For detailed log results see “2-Hour performance test results” on page 26.

#### Individual server metrics

Table 1 shows the sum of I/O’s and the average latency across all storage groups on a per server basis.

***Table 1: Individual Server Performance***

| | **Node 1** | **Node 2** | **Node 3** |
|---|---|---|---|
| **Database I/O** | | | |
| Database Disks Transfers/sec | 1178.01 | 1108.34 | 1174.13 |
| Database Disks Reads/sec | 733.21 | 753.68 | 798.24 |
| Database Disks Write/sec | 344.79 | 354.66 | 375.88 |
| Average Database Disk Read Latency (ms) | 13.85 | 13.96 | 13.15 |
| Average Database Disk Write Latency (ms) | 10.95 | 10.62 | 10.46 |
| **Transaction Log I/O** | | | |
| Log Disks Writes/sec | 81.67 | 84.11 | 88.98 |
| Average Log Disk Write Latency(ms) | 2.84 | 2.84 | 2.83 |

#### Aggregate performance metrics across all servers

Table 2 shows the sum of I/O’s and the average latency across the 3 primary servers in the solution.

***Table 2: Aggregate Server Performance***

| **Database I/O** | |
|---|---|
| Database Disks Transfers/sec | 3360.48 |
| Database Disks Reads/sec | 2285.14 |
| Database Disks Writes/sec | 1075.34 |
| Average Database Disk Read Latency (ms) | 13.65 |
| Average Database Disk Write Latency (ms) | 10.86 |
| **Transaction Log I/O** | |
| Log Disks Writes/sec | 254.767 |
| Average Log Disk Write Latency (ms) | 2.84 |

### 3.7.3 Database backup/recovery performance

Several of the tests in the ESP test framework are used to measure the sequential read rate of the database files and the recovery/replay performance (playing transaction logs into the database).

The database read-only performance test measures the maximum rate at which databases could be backed up using Microsoft Volume Shadow Copy Service (VSS). Table 3 shows the average read performance for a backing up a single database file and all ten database files on a single node.

***Table 3: Database backup read-only performance results***

| | |
|---|---|
| MB read/sec per database | 51.66 |
| MB read/sec total per node (10 databases) | 516.65 |

For detailed log results see “Database backup test results” on page 29.

### 3.7.4 Transaction log recovery/replay performance

The test is to measure the maximum rate at which the log files can be played against the databases. Table 4 shows the average rate for 500 log files played in a single storage group. Each log file is 1 MB in size.

***Table 4: Transaction Log Recovery/Replay Performance***

| | |
|---|---|
| Average time to play one Log file (sec) | 2.20 |

For detailed log results see “Soft recovery test results” on page 31.

# 4 Conclusion

This document is developed by storage solution providers, and reviewed by Microsoft Exchange Product team. The test results/data presented in this document is based on the tests introduced in the ESRP test framework. Customer should not quote the data directly for his/her pre-deployment verification. It is still necessary to go through the exercises to validate the storage design for a specific customer environment.

ESRP program is not designed to be a benchmarking program; tests are not designed to getting the maximum throughput for a giving solution. Rather, it is focused on producing recommendations from vendors for Exchange application. So the data presented in this document should not be used for direct comparisons among the solutions.

The cluster of 4 Lenovo Converged HX7510 appliances proved more than capable of handling the high IOPs generated by 30,000 mailboxes. Part of the reason for this is because each Lenovo Converged HX7510 appliance uses 3 HBAs for the 24 drives. As a consequence the cluster has an average IOPS rate that is 23% better and an average latency improvement of 2 milliseconds over other systems with less HBAs.

# Appendix A – Test results

This section provides test results from 1 of the 3 primary mailbox servers under test: VM1. All server's test results are comparable to each another.

## 24-Hour stress/reliability test results

Checksum Statistics - All

| Database | Seen pages | Bad pages | Correctable pages | Wrong page-number pages | File length / seconds taken |
|---|---|---|---|---|---|
| **E:\Jetstress001001.edb** | 32215552 | 0 | 0 | 0 | 1006736 MB/27425 sec |
| **F:\Jetstress002001.edb** | 32215552 | 0 | 0 | 0 | 1006736 MB/26811 sec |
| **G:\Jetstress003001.edb** | 32216064 | 0 | 0 | 0 | 1006752 MB/26790 sec |
| **H:\Jetstress004001.edb** | 32216064 | 0 | 0 | 0 | 1006752 MB/26779 sec |
| **I:\Jetstress005001.edb** | 32215552 | 0 | 0 | 0 | 1006736 MB/26763 sec |
| **J:\Jetstress006001.edb** | 32215808 | 0 | 0 | 0 | 1006744 MB/26759 sec |
| **K:\Jetstress007001.edb** | 32216320 | 0 | 0 | 0 | 1006760 MB/26772 sec |
| **L:\Jetstress008001.edb** | 32215552 | 0 | 0 | 0 | 1006736 MB/26782 sec |
| **M:\Jetstress009001.edb** | 32215552 | 0 | 0 | 0 | 1006736 MB/26806 sec |
| **N:\Jetstress010001.edb** | 32215296 | 0 | 0 | 0 | 1006728 MB/26825 sec |
| **(Sum)** | 322157312 | 0 | 0 | 0 | 10067416 MB/27426 sec |

Disk Subsystem Performance (of checksum)

| LogicalDisk | Avg. Disk sec/Read | Avg. Disk sec/Write | Disk Reads/sec | Disk Writes/sec | Avg. Disk Bytes/Read |
|---|---|---|---|---|---|
| **E:** | 0.050 | 0.000 | 584.358 | 0.000 | 65536.000 |
| **F:** | 0.049 | 0.000 | 600.062 | 0.000 | 65536.000 |
| **G:** | 0.049 | 0.000 | 600.847 | 0.000 | 65536.000 |
| **H:** | 0.049 | 0.000 | 601.559 | 0.000 | 65536.000 |
| **I:** | 0.049 | 0.000 | 601.839 | 0.000 | 65536.000 |
| **J:** | 0.049 | 0.000 | 601.982 | 0.000 | 65536.000 |
| **K:** | 0.049 | 0.000 | 601.798 | 0.000 | 65536.000 |
| **L:** | 0.049 | 0.000 | 601.210 | 0.000 | 65536.000 |
| **M:** | 0.049 | 0.000 | 600.056 | 0.000 | 65536.000 |
| **N:** | 0.048 | 0.000 | 600.085 | 0.000 | 65536.000 |

Memory System Performance (of checksum)

| Counter | Average | Minimum | Maximum |
|---|---|---|---|
| **% Processor Time** | 42.603 | 8.608 | 58.771 |
| **Available MBytes** | 201583.298 | 201514.000 | 201642.000 |
| **Free System Page Table Entries** | 16467022.549 | 16466370.000 | 16467748.000 |
| **Transition Pages RePurposed/sec** | 0.000 | 0.000 | 0.000 |
| **Pool Nonpaged Bytes** | 74188775.325 | 74051584.000 | 74526720.000 |
| **Pool Paged Bytes** | 133683740.600 | 133578752.000 | 138612736.000 |

Test Log

8/16/2016 5:05:09 AM -- Preparing for testing ...
8/16/2016 5:05:22 AM -- Attaching databases ...
8/16/2016 5:05:22 AM -- Preparations for testing are complete.
8/16/2016 5:05:22 AM -- Starting transaction dispatch ..
8/16/2016 5:05:22 AM -- Database cache settings: (minimum: 320.0 MB, maximum: 2.5 GB)
8/16/2016 5:05:22 AM -- Database flush thresholds: (start: 25.6 MB, stop: 51.2 MB)
8/16/2016 5:05:33 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 200 msec/read).
8/16/2016 5:05:33 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 200 msec/write).
8/16/2016 5:05:34 AM -- Operation mix: Sessions 22, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
8/16/2016 5:05:34 AM -- Performance logging started (interval: 15000 ms).
8/16/2016 5:05:34 AM -- Attaining prerequisites:
8/16/2016 5:09:41 AM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 2420658000.0 (lower bound: 2415919000.0, upper bound: none)
8/17/2016 5:09:41 AM -- Performance logging has ended.
8/17/2016 5:09:41 AM -- JetInterop batch transaction stats: 217143, 217143, 217143, 217143, 217143, 217143, 217142, 217142, 217142 and 217142.
8/17/2016 5:09:41 AM -- Dispatching transactions ends.
8/17/2016 5:09:42 AM -- Shutting down databases ...
8/17/2016 5:09:45 AM -- Instance2744.1 (complete), Instance2744.2 (complete), Instance2744.3 (complete), Instance2744.4 (complete), Instance2744.5 (complete), Instance2744.6 (complete), Instance2744.7 (complete), Instance2744.8 (complete), Instance2744.9 (complete) and Instance2744.10 (complete)
8/17/2016 5:09:45 AM -- C:\Program Files\Exchange Jetstress\24HR Test\Stress_2016_8_16_5_5_33.blg has 5768 samples.
8/17/2016 5:09:45 AM -- Creating test report ...
8/17/2016 5:11:11 AM -- Instance2744.1 has 17.3 for I/O Database Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.1 has 3.2 for I/O Log Writes Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.1 has 3.2 for I/O Log Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.2 has 14.8 for I/O Database Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.2 has 3.1 for I/O Log Writes Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.2 has 3.1 for I/O Log Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.3 has 14.7 for I/O Database Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.3 has 3.1 for I/O Log Writes Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.3 has 3.1 for I/O Log Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.4 has 14.7 for I/O Database Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.4 has 3.1 for I/O Log Writes Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.4 has 3.1 for I/O Log Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.5 has 14.8 for I/O Database Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.5 has 3.2 for I/O Log Writes Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.5 has 3.2 for I/O Log Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.6 has 14.7 for I/O Database Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.6 has 3.0 for I/O Log Writes Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.6 has 3.0 for I/O Log Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.7 has 14.6 for I/O Database Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.7 has 3.1 for I/O Log Writes Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.7 has 3.1 for I/O Log Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.8 has 14.8 for I/O Database Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.8 has 3.1 for I/O Log Writes Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.8 has 3.1 for I/O Log Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.9 has 14.5 for I/O Database Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.9 has 3.1 for I/O Log Writes Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.9 has 3.1 for I/O Log Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.10 has 14.5 for I/O Database Reads Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.10 has 3.1 for I/O Log Writes Average Latency.
8/17/2016 5:11:11 AM -- Instance2744.10 has 3.1 for I/O Log Reads Average Latency.
8/17/2016 5:11:11 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
8/17/2016 5:11:11 AM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
8/17/2016 5:11:11 AM -- C:\Program Files\Exchange Jetstress\24HR Test\Stress_2016_8_16_5_5_33.xml has 5751 samples queried.
8/17/2016 5:11:11 AM -- C:\Program Files\Exchange Jetstress\24HR Test\Stress_2016_8_16_5_5_33.html was saved.
8/17/2016 5:11:12 AM -- Performance logging started (interval: 30000 ms).
8/17/2016 5:11:12 AM -- Verifying database checksums ...
8/17/2016 12:48:18 PM -- E: (100% processed), F: (100% processed), G: (100% processed), H: (100% processed), I: (100% processed), J: (100% processed), K: (100% processed), L: (100% processed), M: (100% processed) and N: (100% processed)
8/17/2016 12:48:19 PM -- Performance logging has ended.
8/17/2016 12:48:19 PM -- C:\Program Files\Exchange Jetstress\24HR Test\DBChecksum_2016_8_17_5_11_11.blg has 913 samples.

# 2-Hour performance test results

Test Summary

| | |
|---|---|
| **Overall Test Result** | **Pass** |
| **Machine Name** | HVM-1 |
| **Test Description** | |
| **Test Start Time** | 8/15/2016 8:22:27 PM |
| **Test End Time** | 8/15/2016 10:26:23 PM |
| **Collection Start Time** | 8/15/2016 8:26:18 PM |
| **Collection End Time** | 8/15/2016 10:26:17 PM |
| **Jetstress Version** | 15.01.0318.000 |
| **ESE Version** | 15.00.0847.030 |
| **Operating System** | Windows Server 2012 R2 Datacenter (6.2.9200.0) |
| **Performance Log** | C:\Program Files\Exchange Jetstress\2HR Test WJ\Performance_2016_8_15_20_22_51.blg |

Database Sizing and Throughput

| | |
|---|---|
| **Achieved Transactional I/O per Second** | 1078.011 |
| **Target Transactional I/O per Second** | 600 |
| **Initial Database Size (bytes)** | 10523340963840 |
| **Final Database Size (bytes)** | 10526100815872 |
| **Database Files (Count)** | 10 |

Jetstress System Parameters

| | |
|---|---|
| **Thread Count** | 22 |
| **Minimum Database Cache** | 320.0 MB |
| **Maximum Database Cache** | 2560.0 MB |
| **Insert Operations** | 40% |
| **Delete Operations** | 20% |
| **Replace Operations** | 5% |
| **Read Operations** | 35% |
| **Lazy Commits** | 70% |
| **Run Background Database Maintenance** | True |
| **Number of Copies per Database** | 2 |

Database Configuration

| | |
|---|---|
| **Instance2744.1** | Log path: E:\Logs<br>Database: E:\Jetstress001001.edb |
| **Instance2744.2** | Log path: F:\Logs<br>Database: F:\Jetstress002001.edb |
| **Instance2744.3** | Log path: G:\Logs<br>Database: G:\Jetstress003001.edb |
| **Instance2744.4** | Log path: H:\Logs<br>Database: H:\Jetstress004001.edb |
| **Instance2744.5** | Log path: I:\Logs<br>Database: I:\Jetstress005001.edb |
| **Instance2744.6** | Log path: J:\Logs<br>Database: J:\Jetstress006001.edb |
| **Instance2744.7** | Log path: K:\Logs<br>Database: K:\Jetstress007001.edb |
| **Instance2744.8** | Log path: L:\Logs<br>Database: L:\Jetstress008001.edb |
| **Instance2744.9** | Log path: M:\Logs<br>Database: M:\Jetstress009001.edb |
| **Instance2744.10** | Log path: N:\Logs<br>Database: N:\Jetstress010001.edb |

Transactional I/O Performance

| MSExchange Database ==> Instances | I/O Database Reads Average Latency (msec) | I/O Database Writes Average Latency (msec) | I/O Database Reads/sec | I/O Database Writes/sec | I/O Database Reads Average Bytes | I/O Database Writes Average Bytes | I/O Log Reads Average Latency (msec) | I/O Log Writes Average Latency (msec) | I/O Log Reads/sec | I/O Log Writes/sec | I/O Log Reads Average Bytes | I/O Log Writes Average Bytes |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| **Instance2744.1** | 16.382 | 11.963 | 73.464 | 34.860 | 34301.360 | 34717.476 | 0.000 | 2.891 | 0.000 | 8.253 | 0.000 | 20303.060 |
| **Instance2744.2** | 13.579 | 12.112 | 73.166 | 34.562 | 34220.772 | 34745.901 | 0.000 | 2.883 | 0.000 | 8.273 | 0.000 | 20325.063 |
| **Instance2744.3** | 13.570 | 11.876 | 73.220 | 34.487 | 34250.416 | 34714.313 | 0.000 | 2.911 | 0.000 | 8.152 | 0.000 | 20373.865 |
| **Instance2744.4** | 13.507 | 11.765 | 73.207 | 34.252 | 34327.959 | 34772.280 | 0.000 | 2.912 | 0.000 | 8.083 | 0.000 | 20662.064 |
| **Instance2744.5** | 14.025 | 11.271 | 73.347 | 34.436 | 34299.922 | 34716.872 | 0.000 | 2.823 | 0.000 | 8.156 | 0.000 | 20331.077 |
| **Instance2744.6** | 13.921 | 10.882 | 73.351 | 34.356 | 34330.468 | 34735.474 | 0.000 | 2.793 | 0.000 | 8.115 | 0.000 | 20402.288 |
| **Instance2744.7** | 13.405 | 10.626 | 73.357 | 34.624 | 34299.104 | 34718.882 | 0.000 | 2.787 | 0.000 | 8.263 | 0.000 | 20280.945 |
| **Instance2744.8** | 13.508 | 10.058 | 73.440 | 34.365 | 34288.856 | 34705.504 | 0.000 | 2.817 | 0.000 | 8.079 | 0.000 | 20119.445 |
| **Instance2744.9** | 13.252 | 9.685 | 73.170 | 34.300 | 34270.735 | 34783.613 | 0.000 | 2.791 | 0.000 | 8.116 | 0.000 | 20554.387 |
| **Instance2744.10** | 13.394 | 9.356 | 73.497 | 34.549 | 34297.920 | 34740.279 | 0.000 | 2.882 | 0.000 | 8.187 | 0.000 | 20271.195 |

Background Database Maintenance I/O Performance

| MSExchange Database ==> Instances | Database Maintenance IO Reads/sec | Database Maintenance IO Reads Average Bytes |
|---|---|---|
| **Instance2744.1** | 9.000 | 261703.339 |
| **Instance2744.2** | 8.995 | 261825.840 |
| **Instance2744.3** | 9.015 | 261751.674 |
| **Instance2744.4** | 8.995 | 261629.449 |
| **Instance2744.5** | 9.009 | 261718.410 |
| **Instance2744.6** | 8.995 | 261708.004 |
| **Instance2744.7** | 9.008 | 261761.733 |
| **Instance2744.8** | 9.022 | 261774.614 |
| **Instance2744.9** | 9.004 | 261697.915 |
| **Instance2744.10** | 9.007 | 261694.352 |

Log Replication I/O Performance

| MSExchange Database ==> Instances | I/O Log Reads/sec | I/O Log Reads Average Bytes |
|---|---|---|
| **Instance2744.1** | 0.714 | 218950.979 |
| **Instance2744.2** | 0.718 | 221383.147 |
| **Instance2744.3** | 0.709 | 220428.474 |
| **Instance2744.4** | 0.712 | 219452.868 |
| **Instance2744.5** | 0.708 | 224308.019 |
| **Instance2744.6** | 0.703 | 219979.442 |
| **Instance2744.7** | 0.713 | 218481.837 |
| **Instance2744.8** | 0.697 | 219880.030 |
| **Instance2744.9** | 0.713 | 218481.837 |
| **Instance2744.10** | 0.708 | 220329.645 |

Total I/O Performance

| MSExchange Database ==> Instances | I/O Database Reads Average Latency (msec) | I/O Database Writes Average Latency (msec) | I/O Database Reads/sec | I/O Database Writes/sec | I/O Database Reads Average Bytes | I/O Database Writes Average Bytes | I/O Log Reads Average Latency (msec) | I/O Log Writes Average Latency (msec) | I/O Log Reads/sec | I/O Log Writes/sec | I/O Log Reads Average Bytes | I/O Log Writes Average Bytes |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| **Instance2744.1** | 16.382 | 11.963 | 82.464 | 34.860 | 59119.842 | 34717.476 | 3.150 | 2.891 | 0.714 | 8.253 | 218950.979 | 20303.060 |
| **Instance2744.2** | 13.579 | 12.112 | 82.161 | 34.562 | 59139.029 | 34745.901 | 3.014 | 2.883 | 0.718 | 8.273 | 221383.147 | 20325.063 |
| **Instance2744.3** | 13.570 | 11.876 | 82.236 | 34.487 | 59191.177 | 34714.313 | 3.017 | 2.911 | 0.709 | 8.152 | 220428.474 | 20373.865 |
| **Instance2744.4** | 13.507 | 11.765 | 82.202 | 34.252 | 59201.161 | 34772.280 | 3.029 | 2.912 | 0.712 | 8.083 | 219452.868 | 20662.064 |
| **Instance2744.5** | 14.025 | 11.271 | 82.356 | 34.436 | 59176.319 | 34716.872 | 3.200 | 2.823 | 0.708 | 8.156 | 224308.019 | 20331.077 |
| **Instance2744.6** | 13.921 | 10.882 | 82.346 | 34.356 | 59168.973 | 34735.474 | 3.005 | 2.793 | 0.703 | 8.115 | 219979.442 | 20402.288 |
| **Instance2744.7** | 13.405 | 10.626 | 82.365 | 34.624 | 59175.802 | 34718.882 | 2.913 | 2.787 | 0.713 | 8.263 | 218481.837 | 20280.945 |
| **Instance2744.8** | 13.508 | 10.058 | 82.463 | 34.365 | 59177.993 | 34705.504 | 2.901 | 2.817 | 0.697 | 8.079 | 219880.030 | 20119.445 |
| **Instance2744.9** | 13.252 | 9.685 | 82.174 | 34.300 | 59190.031 | 34783.613 | 2.933 | 2.791 | 0.713 | 8.116 | 218481.837 | 20554.387 |
| **Instance2744.10** | 13.394 | 9.356 | 82.504 | 34.549 | 59122.611 | 34740.279 | 2.991 | 2.882 | 0.708 | 8.187 | 220329.645 | 20271.195 |

Host System Performance

| Counter | Average | Minimum | Maximum |
|---|---|---|---|
| **% Processor Time** | 10.785 | 2.810 | 18.107 |
| **Available MBytes** | 198926.113 | 198910.000 | 199098.000 |
| **Free System Page Table Entries** | 16467289.208 | 16467082.000 | 16467355.000 |
| **Transition Pages RePurposed/sec** | 0.000 | 0.000 | 0.000 |
| **Pool Nonpaged Bytes** | 73433582.933 | 73031680.000 | 73617408.000 |
| **Pool Paged Bytes** | 127524838.400 | 127504384.000 | 127631360.000 |
| **Database Page Fault Stalls/sec** | 0.000 | 0.000 | 0.000 |

Test Log

```
8/15/2016 8:22:27 PM -- Preparing for testing ...
8/15/2016 8:22:40 PM -- Attaching databases ...
8/15/2016 8:22:40 PM -- Preparations for testing are complete.
8/15/2016 8:22:40 PM -- Starting transaction dispatch ..
8/15/2016 8:22:40 PM -- Database cache settings: (minimum: 320.0 MB, maximum: 2.5 GB)
8/15/2016 8:22:40 PM -- Database flush thresholds: (start: 25.6 MB, stop: 51.2 MB)
8/15/2016 8:22:51 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
8/15/2016 8:22:51 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
8/15/2016 8:22:52 PM -- Operation mix: Sessions 22, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
8/15/2016 8:22:52 PM -- Performance logging started (interval: 15000 ms).
8/15/2016 8:22:52 PM -- Attaining prerequisites:
8/15/2016 8:26:18 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 2421948000.0 (lower bound: 2415919000.0, upper bound: none)
8/15/2016 10:26:19 PM -- Performance logging has ended.
8/15/2016 10:26:19 PM -- JetInterop batch transaction stats: 19655, 19655, 19655, 19655, 19655, 19655, 19655, 19655, 19655 and 19654.
8/15/2016 10:26:19 PM -- Dispatching transactions ends.
8/15/2016 10:26:19 PM -- Shutting down databases ...
8/15/2016 10:26:23 PM -- Instance2744.1 (complete), Instance2744.2 (complete), Instance2744.3 (complete), Instance2744.4 (complete), Instance2744.5
(complete), Instance2744.6 (complete), Instance2744.7 (complete), Instance2744.8 (complete), Instance2744.9 (complete) and Instance2744.10 (complete)
8/15/2016 10:26:23 PM -- C:\Program Files\Exchange Jetstress\2HR Test WJ\Performance_2016_8_15_20_22_51.blg has 493 samples.
8/15/2016 10:26:23 PM -- Creating test report ...
8/15/2016 10:26:26 PM -- Instance2744.1 has 16.4 for I/O Database Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.1 has 2.9 for I/O Log Writes Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.1 has 2.9 for I/O Log Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.2 has 13.6 for I/O Database Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.2 has 2.9 for I/O Log Writes Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.2 has 2.9 for I/O Log Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.3 has 13.6 for I/O Database Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.3 has 2.9 for I/O Log Writes Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.3 has 2.9 for I/O Log Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.4 has 13.5 for I/O Database Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.4 has 2.9 for I/O Log Writes Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.4 has 2.9 for I/O Log Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.5 has 14.0 for I/O Database Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.5 has 2.8 for I/O Log Writes Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.5 has 2.8 for I/O Log Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.6 has 13.9 for I/O Database Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.6 has 2.8 for I/O Log Writes Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.6 has 2.8 for I/O Log Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.7 has 13.4 for I/O Database Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.7 has 2.8 for I/O Log Writes Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.7 has 2.8 for I/O Log Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.8 has 13.5 for I/O Database Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.8 has 2.8 for I/O Log Writes Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.8 has 2.8 for I/O Log Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.9 has 13.3 for I/O Database Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.9 has 2.8 for I/O Log Writes Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.9 has 2.8 for I/O Log Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.10 has 13.4 for I/O Database Reads Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.10 has 2.9 for I/O Log Writes Average Latency.
8/15/2016 10:26:26 PM -- Instance2744.10 has 2.9 for I/O Log Reads Average Latency.
8/15/2016 10:26:26 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
8/15/2016 10:26:26 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
8/15/2016 10:26:26 PM -- C:\Program Files\Exchange Jetstress\2HR Test WJ\Performance_2016_8_15_20_22_51.xml has 479 samples queried.
```

# Database backup test results

Database Backup Statistics - All

| Database Instance | Database Size (MBytes) | Elapsed Backup Time | MBytes Transferred/sec |
|---|---|---|---|
| **Instance2744.1** | 1006728.03 | 07:52:16 | 35.53 |
| **Instance2744.2** | 1006728.03 | 07:42:36 | 36.27 |
| **Instance2744.3** | 1006744.03 | 07:42:43 | 36.26 |
| **Instance2744.4** | 1006744.03 | 07:42:36 | 36.27 |
| **Instance2744.5** | 1006728.03 | 07:42:04 | 36.31 |
| **Instance2744.6** | 1006736.03 | 07:42:11 | 36.30 |
| **Instance2744.7** | 1006752.03 | 07:42:37 | 36.27 |
| **Instance2744.8** | 1006728.03 | 07:42:45 | 36.26 |
| **Instance2744.9** | 1006728.03 | 07:42:37 | 36.27 |
| **Instance2744.10** | 1006720.03 | 07:42:36 | 36.27 |
| **Avg** | | | 36.20 |
| **Sum** | | | 362.01 |

Jetstress System Parameters

| | |
|---|---|
| **Thread Count** | 22 |
| **Minimum Database Cache** | 320.0 MB |
| **Maximum Database Cache** | 2560.0 MB |
| **Insert Operations** | 40% |
| **Delete Operations** | 20% |
| **Replace Operations** | 5% |
| **Read Operations** | 35% |
| **Lazy Commits** | 70% |

Database Configuration

**Instance2744.1** Log path: E:\Logs
Database: E:\Jetstress001001.edb

**Instance2744.2** Log path: F:\Logs
Database: F:\Jetstress002001.edb

**Instance2744.3** Log path: G:\Logs
Database: G:\Jetstress003001.edb

**Instance2744.4** Log path: H:\Logs
Database: H:\Jetstress004001.edb

**Instance2744.5** Log path: I:\Logs
Database: I:\Jetstress005001.edb

**Instance2744.6** Log path: J:\Logs
Database: J:\Jetstress006001.edb

**Instance2744.7** Log path: K:\Logs
Database: K:\Jetstress007001.edb

**Instance2744.8** Log path: L:\Logs
Database: L:\Jetstress008001.edb

**Instance2744.9** Log path: M:\Logs
Database: M:\Jetstress009001.edb

**Instance2744.10** Log path: N:\Logs
Database: N:\Jetstress010001.edb

Transactional I/O Performance

| MSExchange Database ==> Instances | I/O Database Reads Average Latency (msec) | I/O Database Writes Average Latency (msec) | I/O Database Reads/sec | I/O Database Writes/sec | I/O Database Reads Average Bytes | I/O Database Writes Average Bytes | I/O Log Reads Average Latency (msec) | I/O Log Writes Average Latency (msec) | I/O Log Reads/sec | I/O Log Writes/sec | I/O Log Reads Average Bytes | I/O Log Writes Average Bytes |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| **Instance2744.1** | 13.051 | 0.000 | 141.724 | 0.000 | 262144.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| **Instance2744.2** | 12.905 | 0.000 | 145.004 | 0.000 | 262144.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| **Instance2744.3** | 12.919 | 0.000 | 144.812 | 0.000 | 262144.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| **Instance2744.4** | 12.905 | 0.000 | 145.007 | 0.000 | 262144.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| **Instance2744.5** | 12.897 | 0.000 | 145.184 | 0.000 | 262144.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| **Instance2744.6** | 12.899 | 0.000 | 145.125 | 0.000 | 262144.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| **Instance2744.7** | 12.907 | 0.000 | 144.990 | 0.000 | 262144.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| **Instance2744.8** | 12.921 | 0.000 | 144.764 | 0.000 | 262144.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| **Instance2744.9** | 12.907 | 0.000 | 144.999 | 0.000 | 262144.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| **Instance2744.10** | 12.905 | 0.000 | 145.003 | 0.000 | 262144.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |

Host System Performance

| Counter | Average | Minimum | Maximum |
|---|---|---|---|
| **% Processor Time** | 31.517 | 5.027 | 40.812 |
| **Available MBytes** | 201609.912 | 201588.000 | 201614.000 |
| **Free System Page Table Entries** | 16467773.196 | 16467404.000 | 16467872.000 |
| **Transition Pages RePurposed/sec** | 0.000 | 0.000 | 0.000 |
| **Pool Nonpaged Bytes** | 73616132.339 | 73572352.000 | 73703424.000 |
| **Pool Paged Bytes** | 133731284.610 | 133718016.000 | 133869568.000 |
| **Database Page Fault Stalls/sec** | 0.000 | 0.000 | 0.000 |

Test Log

8/17/2016 1:14:40 PM -- Preparing for testing ...
8/17/2016 1:14:53 PM -- Attaching databases ...
8/17/2016 1:14:53 PM -- Preparations for testing are complete.
8/17/2016 1:15:03 PM -- Performance logging started (interval: 30000 ms).
8/17/2016 1:15:03 PM -- Backing up databases ...
8/17/2016 9:07:20 PM -- Performance logging has ended.
8/17/2016 9:07:20 PM -- Instance2744.1 (100% processed), Instance2744.2 (100% processed), Instance2744.3 (100% processed), Instance2744.4 (100% processed), Instance2744.5 (100% processed), Instance2744.6 (100% processed), Instance2744.7 (100% processed), Instance2744.8 (100% processed), Instance2744.9 (100% processed) and Instance2744.10 (100% processed)
8/17/2016 9:07:20 PM -- C:\Program Files\Exchange Jetstress\DB Backup\DatabaseBackup_2016_8_17_13_14_53.blg has 944 samples.
8/17/2016 9:07:20 PM -- Creating test report ...

# Soft recovery test results

Soft-Recovery Statistics - All

| Database Instance | Log files replayed | Elapsed seconds |
|---|---|---|
| **Instance2744.1** | 505 | 1121.8926619 |
| **Instance2744.2** | 507 | 1034.9832663 |
| **Instance2744.3** | 504 | 1027.2257387 |
| **Instance2744.4** | 504 | 1037.0420222 |
| **Instance2744.5** | 511 | 1038.8198341 |
| **Instance2744.6** | 506 | 1056.888243 |
| **Instance2744.7** | 506 | 1056.1317836 |
| **Instance2744.8** | 501 | 1057.1445261 |
| **Instance2744.9** | 507 | 1061.0030249 |
| **Instance2744.10** | 503 | 1055.86615 |
| **Avg** | 505 | 1054.7 |
| **Sum** | 5054 | 10546.9972508 |

Database Configuration

**Instance2744.1** Log path: E:\Logs
Database: E:\Jetstress001001.edb

**Instance2744.2** Log path: F:\Logs
Database: F:\Jetstress002001.edb

**Instance2744.3** Log path: G:\Logs
Database: G:\Jetstress003001.edb

**Instance2744.4** Log path: H:\Logs
Database: H:\Jetstress004001.edb

**Instance2744.5** Log path: I:\Logs
Database: I:\Jetstress005001.edb

**Instance2744.6** Log path: J:\Logs
Database: J:\Jetstress006001.edb

**Instance2744.7** Log path: K:\Logs
Database: K:\Jetstress007001.edb

**Instance2744.8** Log path: L:\Logs
Database: L:\Jetstress008001.edb

**Instance2744.9** Log path: M:\Logs
Database: M:\Jetstress009001.edb

**Instance2744.10** Log path: N:\Logs
Database: N:\Jetstress010001.edb

Transactional I/O Performance

| MSExchange Database ==> Instances | I/O Database Reads Average Latency (msec) | I/O Database Writes Average Latency (msec) | I/O Database Reads/sec | I/O Database Writes/sec | I/O Database Reads Average Bytes | I/O Database Writes Average Bytes | I/O Log Reads Average Latency (msec) | I/O Log Writes Average Latency (msec) | I/O Log Reads/sec | I/O Log Writes/sec | I/O Log Reads Average Bytes | I/O Log Writes Average Bytes |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| **Instance2744.1** | 167.748 | 29.795 | 266.032 | 1.767 | 37661.072 | 29562.435 | 3.303 | 0.000 | 2.209 | 0.000 | 186267.579 | 0.000 |
| **Instance2744.2** | 78.894 | 2.586 | 290.566 | 1.956 | 37881.952 | 29812.455 | 3.121 | 0.000 | 2.445 | 0.000 | 189975.844 | 0.000 |
| **Instance2744.3** | 63.479 | 2.495 | 294.122 | 1.963 | 37902.490 | 29789.091 | 3.015 | 0.000 | 2.454 | 0.000 | 190650.182 | 0.000 |
| **Instance2744.4** | 85.228 | 15.544 | 286.766 | 1.925 | 38215.907 | 29683.953 | 2.838 | 0.000 | 2.406 | 0.000 | 188292.864 | 0.000 |
| **Instance2744.5** | 66.049 | 8.422 | 288.593 | 1.964 | 37853.129 | 30080.000 | 3.122 | 0.000 | 2.455 | 0.000 | 191677.153 | 0.000 |
| **Instance2744.6** | 213.478 | 2.566 | 279.365 | 1.905 | 38195.691 | 29869.292 | 3.027 | 0.000 | 2.382 | 0.000 | 191222.670 | 0.000 |
| **Instance2744.7** | 121.177 | 2.654 | 279.617 | 1.899 | 37931.036 | 29491.200 | 3.065 | 0.000 | 2.374 | 0.000 | 188750.049 | 0.000 |
| **Instance2744.8** | 119.421 | 5.104 | 285.903 | 1.876 | 37842.560 | 29995.323 | 3.373 | 0.000 | 2.345 | 0.000 | 191953.701 | 0.000 |
| **Instance2744.9** | 109.816 | 2.427 | 279.826 | 1.892 | 38061.136 | 29503.755 | 3.264 | 0.000 | 2.365 | 0.000 | 188817.525 | 0.000 |
| **Instance2744.10** | 141.867 | 2.431 | 284.877 | 1.899 | 37831.636 | 29617.231 | 3.007 | 0.000 | 2.374 | 0.000 | 188745.823 | 0.000 |

Background Database Maintenance I/O Performance

| MSExchange Database ==> Instances | Database Maintenance IO Reads/sec | Database Maintenance IO Reads Average Bytes |
|---|---|---|
| **Instance2744.1** | 0.000 | 0.000 |
| **Instance2744.2** | 0.000 | 0.000 |
| **Instance2744.3** | 0.000 | 0.000 |
| **Instance2744.4** | 0.000 | 0.000 |
| **Instance2744.5** | 0.000 | 0.000 |
| **Instance2744.6** | 0.000 | 0.000 |
| **Instance2744.7** | 0.000 | 0.000 |
| **Instance2744.8** | 0.000 | 0.000 |
| **Instance2744.9** | 0.000 | 0.000 |
| **Instance2744.10** | 0.000 | 0.000 |

Total I/O Performance

| MSExchange Database ==> Instances | I/O Database Reads Average Latency (msec) | I/O Database Writes Average Latency (msec) | I/O Database Reads/sec | I/O Database Writes/sec | I/O Database Reads Average Bytes | I/O Database Writes Average Bytes | I/O Log Reads Average Latency (msec) | I/O Log Writes Average Latency (msec) | I/O Log Reads/sec | I/O Log Writes/sec | I/O Log Reads Average Bytes | I/O Log Writes Average Bytes |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| **Instance2744.1** | 167.748 | 29.795 | 266.032 | 1.767 | 37661.072 | 29562.435 | 3.303 | 0.000 | 2.209 | 0.000 | 186267.579 | 0.000 |
| **Instance2744.2** | 78.894 | 2.586 | 290.566 | 1.956 | 37881.952 | 29812.455 | 3.121 | 0.000 | 2.445 | 0.000 | 189975.844 | 0.000 |
| **Instance2744.3** | 63.479 | 2.495 | 294.122 | 1.963 | 37902.490 | 29789.091 | 3.015 | 0.000 | 2.454 | 0.000 | 190650.182 | 0.000 |
| **Instance2744.4** | 85.228 | 15.544 | 286.766 | 1.925 | 38215.907 | 29683.953 | 2.838 | 0.000 | 2.406 | 0.000 | 188292.864 | 0.000 |
| **Instance2744.5** | 66.049 | 8.422 | 288.593 | 1.964 | 37853.129 | 30080.000 | 3.122 | 0.000 | 2.455 | 0.000 | 191677.153 | 0.000 |
| **Instance2744.6** | 213.478 | 2.566 | 279.365 | 1.905 | 38195.691 | 29869.292 | 3.027 | 0.000 | 2.382 | 0.000 | 191222.670 | 0.000 |
| **Instance2744.7** | 121.177 | 2.654 | 279.617 | 1.899 | 37931.036 | 29491.200 | 3.065 | 0.000 | 2.374 | 0.000 | 188750.049 | 0.000 |
| **Instance2744.8** | 119.421 | 5.104 | 285.903 | 1.876 | 37842.560 | 29995.323 | 3.373 | 0.000 | 2.345 | 0.000 | 191953.701 | 0.000 |
| **Instance2744.9** | 109.816 | 2.427 | 279.826 | 1.892 | 38061.136 | 29503.755 | 3.264 | 0.000 | 2.365 | 0.000 | 188817.525 | 0.000 |
| **Instance2744.10** | 141.867 | 2.431 | 284.877 | 1.899 | 37831.636 | 29617.231 | 3.007 | 0.000 | 2.374 | 0.000 | 188745.823 | 0.000 |

Host System Performance

| Counter | Average | Minimum | Maximum |
|---|---|---|---|
| **% Processor Time** | 24.248 | 0.000 | 53.647 |
| **Available MBytes** | 199023.375 | 198886.000 | 201498.000 |
| **Free System Page Table Entries** | 16467663.300 | 16467338.000 | 16467790.000 |
| **Transition Pages RePurposed/sec** | 0.000 | 0.000 | 0.000 |
| **Pool Nonpaged Bytes** | 74128727.798 | 73629696.000 | 74764288.000 |
| **Pool Paged Bytes** | 134608356.274 | 134586368.000 | 134643712.000 |
| **Database Page Fault Stalls/sec** | 0.000 | 0.000 | 0.000 |

Test Log

```
8/17/2016 10:48:50 PM -- Preparing for testing ...
8/17/2016 10:49:02 PM -- Attaching databases ...
8/17/2016 10:49:02 PM -- Preparations for testing are complete.
8/17/2016 10:49:02 PM -- Starting transaction dispatch ..
8/17/2016 10:49:02 PM -- Database cache settings: (minimum: 320.0 MB, maximum: 2.5 GB)
8/17/2016 10:49:02 PM -- Database flush thresholds: (start: 25.6 MB, stop: 51.2 MB)
8/17/2016 10:49:13 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
8/17/2016 10:49:13 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
8/17/2016 10:49:14 PM -- Operation mix: Sessions 22, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
8/17/2016 10:49:14 PM -- Performance logging started (interval: 15000 ms).
8/17/2016 10:49:14 PM -- Generating log files ...
8/18/2016 12:44:51 AM -- E:\Logs (101.0% generated), F:\Logs (101.4% generated), G:\Logs (100.8% generated), H:\Logs (100.6% generated), I:\Logs
(102.2% generated), J:\Logs (101.2% generated), K:\Logs (101.2% generated), L:\Logs (100.2% generated), M:\Logs (101.4% generated) and N:\Logs
(100.6% generated)
8/18/2016 12:44:51 AM -- Performance logging has ended.
8/18/2016 12:44:51 AM -- JetInterop batch transaction stats: 17072, 17072, 17072, 17071, 17071, 17071, 17071, 17071, 17071 and 17071.
8/18/2016 12:44:51 AM -- Dispatching transactions ends.
8/18/2016 12:44:52 AM -- Shutting down databases ...
8/18/2016 12:45:00 AM -- Instance2744.1 (complete), Instance2744.2 (complete), Instance2744.3 (complete), Instance2744.4 (complete), Instance2744.5
(complete), Instance2744.6 (complete), Instance2744.7 (complete), Instance2744.8 (complete), Instance2744.9 (complete) and Instance2744.10 (complete)
8/18/2016 12:45:00 AM -- C:\Program Files\Exchange Jetstress\Soft recovery\Performance_2016_8_17_22_49_13.blg has 462 samples.
8/18/2016 12:45:00 AM -- Creating test report ...
8/18/2016 12:45:02 AM -- Instance2744.1 has 18.5 for I/O Database Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.1 has 3.4 for I/O Log Writes Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.1 has 3.4 for I/O Log Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.2 has 15.5 for I/O Database Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.2 has 3.0 for I/O Log Writes Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.2 has 3.0 for I/O Log Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.3 has 15.4 for I/O Database Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.3 has 3.1 for I/O Log Writes Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.3 has 3.1 for I/O Log Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.4 has 15.5 for I/O Database Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.4 has 3.1 for I/O Log Writes Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.4 has 3.1 for I/O Log Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.5 has 15.6 for I/O Database Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.5 has 3.1 for I/O Log Writes Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.5 has 3.1 for I/O Log Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.6 has 15.7 for I/O Database Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.6 has 3.1 for I/O Log Writes Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.6 has 3.1 for I/O Log Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.7 has 15.6 for I/O Database Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.7 has 3.1 for I/O Log Writes Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.7 has 3.1 for I/O Log Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.8 has 15.7 for I/O Database Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.8 has 3.0 for I/O Log Writes Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.8 has 3.0 for I/O Log Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.9 has 15.5 for I/O Database Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.9 has 3.0 for I/O Log Writes Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.9 has 3.0 for I/O Log Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.10 has 15.6 for I/O Database Reads Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.10 has 3.1 for I/O Log Writes Average Latency.
8/18/2016 12:45:02 AM -- Instance2744.10 has 3.1 for I/O Log Reads Average Latency.
8/18/2016 12:45:02 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
8/18/2016 12:45:02 AM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
8/18/2016 12:45:02 AM -- C:\Program Files\Exchange Jetstress\Soft recovery\Performance_2016_8_17_22_49_13.xml has 461 samples queried.
8/18/2016 12:45:02 AM -- C:\Program Files\Exchange Jetstress\Soft recovery\Performance_2016_8_17_22_49_13.html was saved.
8/18/2016 12:45:02 AM -- Performance logging started (interval: 4000 ms).
8/18/2016 12:45:02 AM -- Recovering databases ...
8/18/2016 1:03:45 AM -- Performance logging has ended.
```

# Appendix B - Resources

- Lenovo Converged HX Series Product Guide
  lenovopress.com/lp0059
- Lenovo Converged HX Series landing page
  shop.lenovo.com/us/en/systems/converged-systems/hx_series
- Nutanix Portal (requires registration)
  portal.nutanix.com
- Nutanix Bible
  nutanixbible.com/
- Nutanix Tech Note: VMware vSphere Networking on Nutanix
  go.nutanix.com/rs/nutanix/images/Nutanix_TechNote-VMware_vSphere_Networking_with_Nutanix.pdf

# Trademarks and special notices



References in this document to Lenovo products or services do not imply that Lenovo intends to make them available in every country.

Lenovo, the Lenovo logo, ThinkCentre, ThinkVision, ThinkVantage, ThinkPlus and Rescue and Recovery are trademarks of Lenovo.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Nutanix is a trademark of Nutanix, Inc. in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used Lenovo products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-Lenovo products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by Lenovo. Sources for non-Lenovo list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. Lenovo has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-Lenovo products. Questions on the capability of non-Lenovo products should be addressed to the supplier of those products.

All statements regarding Lenovo future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local Lenovo office or Lenovo authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in Lenovo product announcements. The information is presented here to communicate Lenovo’s current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard Lenovo benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-Lenovo websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this Lenovo product and use of those websites is at your own risk.